arXiv:1511.06627v1 [cs.CV] 20 Nov 2015

Towards Arbitrary-View Face Alignment by Recommendation Trees*

Shizhan Zhu[1] Cheng Li[2] Chen Change Loy[1,3] Xiaoou Tang[1,3]
[1]Department of Information Engineering, The Chinese University of Hong Kong
[2]SenseTime Group
[3]Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences
zs014@ie.cuhk.edu.hk, chengli@sensetime.com, ccloy@ie.cuhk.edu.hk, xtang@ie.cuhk.edu.hk

Abstract

Learning to simultaneously handle face alignment of arbitrary views, e.g. *frontal and profile views, appears to be more challenging than we thought. The difficulties lay in i) accommodating the complex appearance-shape relations exhibited in different views, and ii) encompassing the varying landmark point sets due to self-occlusion and different landmark protocols. Most existing studies approach this problem via training multiple viewpoint-specific models, and conduct head pose estimation for model selection. This solution is intuitive but the performance is highly susceptible to inaccurate head pose estimation. In this study, we address this shortcoming through learning an Ensemble of Model Recommendation Trees (EMRT), which is capable of selecting optimal model configuration without prior head pose estimation. The unified framework seamlessly handles different viewpoints and landmark protocols, and it is trained by optimising directly on landmark locations, thus yielding superior results on arbitrary-view face alignment. This is the first study that performs face alignment on the full AFLW dataset with faces of different views including profile view. State-of-the-art performances are also reported on MultiPIE and AFW datasets containing both frontal- and profile-view faces.*

1. Introduction

Face alignment is essential for many facial analysis tasks such as facial expression detection [11, 28], attribute recognition [9], face image processing [25], face recognition and verification [22, 5, 21, 10, 12]. Significant improvements have been achieved in face alignment with the emergence of discriminative regression methods [24, 17, 13]. Nevertheless, most of these methods tend to confine their solutions on faces with limited view range, in which at least both two eyes are visible and all landmarks can be explicitly annotated. Some studies [31, 23] can handle multi-view face alignment, but they only reported results on faces with $\sim \pm 45°$ viewpoint variation. When view changes dramatically, these models often produce erroneous shape estimation (Fig. 1). However, non-frontal face constitutes a considerable fraction of faces found in various sources, such as web images, social network photos, and surveillance videos. Considering alignment only for near-frontal view severely limits the scalability of many subsequent facial analyses in real-world environment. In this paper, we wish to study the limitations of using a single discriminative regression model [24, 17, 13], and formulate a unified model recommendation framework to handle arbitrary view poses up to profile view $\sim \pm 90°$.

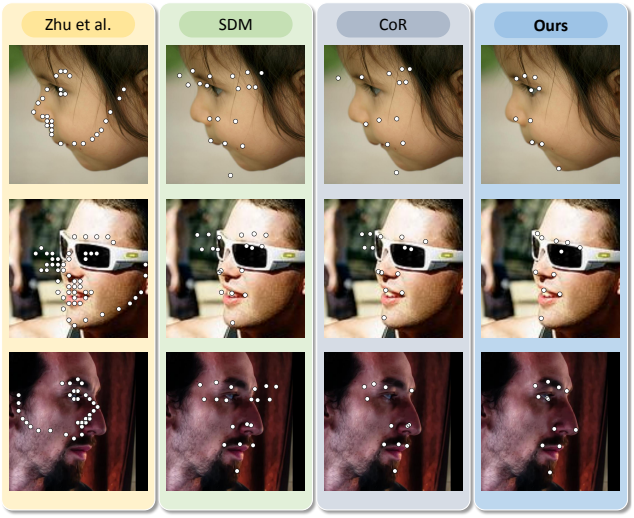

Figure 1. Results on four methods (from left to right): [32] models facial structure by mixture of trees; SDM [24] is a representative discriminative regression method that achieves state-of-the-art performance on frontal-face-biased datasets [1, 15]; CoR [27] is specially proposed to handle external-object occlusion; and our proposed method based on model recommendation.

Simultaneously considering face alignment for arbitrary

*This is our original submission to ICCV 2015.

view poses is non-trivial: *i)* For large view variation some landmarks such as those on half of the face might become invisible due to self-occlusion. Solving this problem requires not only the capability to infer the visibility of occluded points, but also the ability to handle varying landmark point sets. For instance, MultiPIE has 39 points for profile-view and 68 points for frontal-view. *ii)* Importantly, due to the inherently multimodal and complex appearance-shape relations exhibited in multi-view faces, a single model would easily fail to fit the landmarks (Fig. 1). The aforementioned challenges demand multiple face alignment models specially trained on different view poses, so as to handle different landmark protocols and multimodal appearance-shape relations. How to jointly and seamlessly consider all models for aligning an arbitrary face pose remains an understudied problem.

There are a few attempts [8, 30] to address the arbitrary-view face alignment. These methods first estimate the head pose and subsequently choose the corresponding alignment model from a model pool for landmark estimation. Nonetheless, head pose estimation in unconstrained settings itself is a challenging task. Existing methods mostly rely on global features for estimation, which can be susceptible to illumination, occlusion, and background clutter. Error in head pose estimation inevitably affects the performance of subsequent face alignment. Approaches have also been proposed to handle occlusion caused by external objects such as hands, hairs, or sunglasses [3, 27]. Nonetheless, there are intrinsic differences between the handling of external occlusion and self-occlusion caused by view variation. In particular, given a specific pose, the underlying shape constraint remains unchanged in the presence of external occlusion. Even if some facial parts are occluded, accurate landmarks can still be inferred using a single model. On the contrary, self-occlusion caused by view changes is a more challenging problem. Our experiments show that a direct application of face alignment specifically designed for handling external occlusion fails to cope with profile-view poses (CoR [27] in Fig. 1).

Handling face alignment with different view poses requires a framework that can seamlessly handle different face models and landmark protocols from frontal view to profile view, and select optimal model configuration given an arbitrary pose. To this end, we cast the problem as a model recommendation problem and propose a framework called Ensemble of Model Recommendation Trees (EMRT). Decision forests [2] have been extensively used in many applications including facial alignment due to its effectiveness and efficiency. In this study, we reformulate it for the purpose of model recommendation. In the proposed framework, we first prepare a model pool by training a set of alignment models, each of which concentrates on processing specific cluster of face samples with similar face angles and landmark protocols[1]. Given a face image with unknown pose, the EMRT takes features based on the shape responses from different models as input, and produces a model rating vector that encodes the preference on different models. Face alignment is then accomplished by aggregating the consensus of different models based on the rating.

The main contributions of this study are as follows:
(1) We propose a novel approach that is capable of handling arbitrary pose face alignment, up to profile view. To our knowledge, this is the first study that is capable of training and evaluating on the challenging **full** Annotated Facial Landmarks in the Wild (AFLW) dataset [14] [2]. State-of-the-art performances are also achieved on MultiPIE [20] and AFW [32], which contain challenging profile-view faces.
(2) We formulate a new forest-based recommendation framework with appealing properties. Firstly, it has a new splitting function that jointly optimises the partitions of features and the learning of model rating vector. The splitting function is formulated such that it optimises directly the landmark locations, thus leading to accurate alignment given arbitrary poses. Secondly, the objective function is optimised to produce a model rating vector that captures the unique correlation between models so that semantically nearer models (*e.g.* those trained with faces in nearby face angle range) share closer values in the model rating vector, than those further away in the model space. Thanks to these properties, the proposed model recommendation approach is more robust than other model selection alternatives in our task.

Although we mainly focus on face alignment of arbitrary views, our framework also naturally deals with partial occlusion and capable of inferring the visibility of landmarks.

2. Related Work

Many approaches have been proposed for solving the face alignment problem. The classic methods include the Active Appearance Model (AAM) [6] and its variants. AAM models faces by both appearance and shape information and optimises them in a holistic manner. Facial parts detection based method [1], also known as Constrained Local Model, aim to detect facial parts with several part detectors, and use global constraints to estimate reasonable shapes. Discriminative regression methods [4, 24, 17, 13] regress expected shape update based on local appearance information extracted from all current estimated landmarks. Such methods have achieved state-of-the-art accuracy. Our

[1]Note that our framework can naturally handle models encompassing different in-plane and out-of-plane rotations. In addition, the framework flexibly handles varying number of models depending on the complexity of the application scenarios, and it can be used along with different existing discriminative regression methods [24, 17, 13].

[2]The full AFLW dataset contains 25,993 images. Existing studies [7, 19, 26, 29] only select $\leq$25% images with near-to-frontal views for evaluations.

framework allows these state-of-the-art approaches to be applicable to non-frontal cases via recommendation over multiple models. A number of existing studies propose to solve the multi-view face alignment problem [31, 23]. These studies differ significantly from our work since they confine the yaw angle of faces within the range of $\sim \pm 45°$ [31]. In [23], the dataset is chosen from the 300-W dataset [18] with an approximate range of $\sim \pm 45°$, which is far smaller than the head pose range considered in this study.

Our work is closer to a few alignment studies that address the face alignment problem in a multimodal manner:

Conditional Regression Forest [8]. In this work, the facial points are estimated after a rough estimation of head pose. The head pose is estimated according to the global appearance of the face image and regarded as the prior probability of each alignment model. This method works quite well and gains improvement on traditional face alignment benchmarks, *e.g.* LFPW [1] or LFW [8]. When this method is extended to dataset with extremely large pose variations, *e.g.* AFLW, the method fails due to erroneous head pose estimation, no matter how well each face alignment model is trained.

Iterative Multi-Output Random Forests [30]. This study extends [8] by iteratively re-estimating the head pose and facial expression based on the current estimated shape, and then re-estimating the facial landmarks with the updated prior model probability. Further improvement is achieved in this work compared to [8]. Similar to [8], the algorithm begins with head pose estimation based on global appearance. When applied to faces with large pose variation, if the initial head pose estimation is incorrect, the error is rather hard to be rectified during iterative optimisation since the initial shape is estimated with the wrong prior probability.

Consensus of Regressors [27]. The study applies a mixture of models to handle different types of occluded faces. The method works well on occlusion benchmarks like COFW [3]. However, the method yields poor performance on profile-view cases because all regression models assume a limited pose variation. From our experiments, we observe that even if models are provided to handle large pose variations, suboptimal results may be still generated due to the failure in drawing consensus amongst recommended landmarks with large spatial variations.

Model Recommendation. Model recommendation has been proposed for action recognition [16]. The study shows that a recommendation system is important for choosing one model out of a large set of possibilities. Their method follows the conventional collaborative filtering scheme, where user's rating is required on a small subset of models, and the goal is to use that subset of ratings to predict the ratings of remaining models. This is in contrast to our task where user rating is not available.

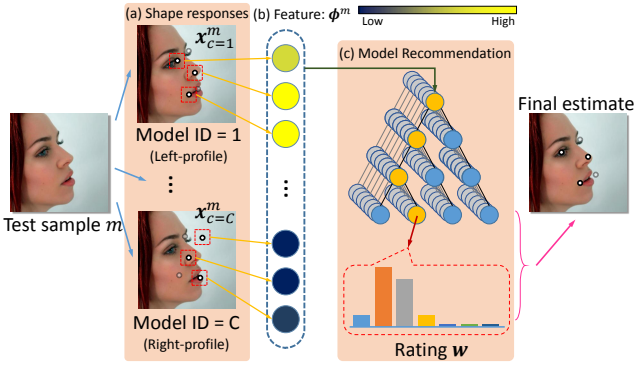

Figure 2. There are three steps to align a face image with arbitrary pose: (a) extracting shape responses from *model pool* (note that the model pool could cover the modeling of various in-plane or out-plane face rotation and external-object occlusion), (b) forming *recommendation feature* vector based on model shape responses, and (c) model recommendation using the proposed *Ensemble of Model Recommendation Trees*.

Category	Notation	Meaning
Index	$c = 1, ..., C$	Index of model
	$m = 1, ..., M$	Index of sample
	$n = 1, ..., N$	Index of facial landmark
	$f = 1, ..., F$	Index of feature
	j	Index of node in a tree
Training Set	$\boldsymbol{x}^m_{c,n} \in \mathcal{X}, \mathbb{R}^2$	Model shape responses
	$\boldsymbol{y}^m_n \in \mathcal{Y}, \mathbb{R}^2$	Shape ground truth
	$\{m, n\} \in \mathcal{V}$	Visibility ground truth
	$\phi^m_f \in \mathcal{F}$	Feature
	$l^m_n \in \mathcal{L}$	Label for classification
Model	$\boldsymbol{b}_c \in \{0, 1\}^N$	Model visibility protocol
Rating	$\boldsymbol{w}_j \in \mathbb{R}^C$	Rating

Table 1. Mathematical notations used in this study.

3. Arbitrary-View Face Alignment

Our aim is to predict the N facial landmark locations given a face image with arbitrary view up to profile view, with potentially in-plane rotation, self- or external-object occlusion. In addition to landmark locations, we also want to estimate a confidence value that indicates the visibility of each landmark. Figure 2 provides an overview of using the proposed framework for aligning a face with arbitrary pose. We explain the training process involved in each step as follows. Table 1 summarises the notations we used in this study.

a) Model pool construction. We construct a *model pool* of discriminative regression alignment models as depicted in Fig. 2(a). We assume there are C models in the pool and M training samples. Each **model** c is trained to accept a facial image and output the predicted landmarks $\{\boldsymbol{x}_{c,n}\}_{n=1}^N$, where $c \in \{1, \ldots, C\}$, and N is the number of landmarks. Here $\boldsymbol{x}$ is the (x, y) coordinate of one landmark. The land-

marks predicted by all models over all samples collectively form the model shapes responses $\mathcal{X} = \{\boldsymbol{x}_{c,n}^m\}_{c=1,n=1,m=1}^{C,\ N,\ M}$.

Each alignment model is trained to handle only a specific cluster of faces with specific view range and landmark visibility protocol[3]. Models are different from each other, *e.g.* there can be a model specific to left-profile, or a model specific to 45° out-of-plane rotated with mouth occluded. We denote landmark protocol of model c by a binary mask vector $\boldsymbol{b}_c \in \{0,1\}^N$, 1 for visible and 0 for invisible (derived from ground truth visibility $\mathcal{V}$). To train each model, we provide it with the labelled facial landmark locations $\mathcal{Y} = \{\boldsymbol{y}_n^m\}_{\{m,n\}\in\mathcal{V}}$ and visibility $\mathcal{V}$. Here $\{m,n\} \in \mathcal{V}$ only when the n-th landmarks in the m-th face is visible. In other words, the training samples for a model must contain all visible landmarks defined by the associated model landmark protocol. Note that we still force all the models to output the full set of landmarks, including those defined as invisible by the model, so as to facilitate the following recommendation process, *i.e.* $\boldsymbol{x}_{c,n}^m \in \mathcal{X}$ even if $(\boldsymbol{b}_c)_n = 0$.

Consequently, each alignment model is capable of handling its specific type of faces easily; the hard part lies in the way to infer the optimal model configuration given a face with potentially arbitrary pose and landmarks visibility.

b) Recommendation features extraction. All alignment models in the pool collectively produce the model shape responses $\mathcal{X}$ given a face image, as depicted in Fig. 2(a). From the given model shape responses we extract recommendation features as input to the EMRT. Formally, we represent the recommendation feature of the m-th face image as $\boldsymbol{\phi}^m = (\phi_1^m, \ldots, \phi_F^m)^\mathsf{T}$, where each dimension a facial part detection score of a visible landmark, and the dimension is denoted by $F = \sum_{c=1}^{C} \|\boldsymbol{b}_c\|_1$. To obtain the part detection score, we apply a facial part detector (SIFT+SVM) on all visible landmarks defined in all models in the pool, and use the resulting detection scores to form the recommendation feature of a face. Such feature is more discriminative than global appearance features used in the prior-head-pose-estimation method. We collectively denote all features of M training images as $\mathcal{F} = \{\phi_f^m\}_{m=1,f=1}^{M,\ F}$.

c) Model Recommendation by ensemble of recommendation trees. The core purpose of a recommendation system is to estimate a rating that describe the fitness of a model on a given task. We denote a model rating vector as $\boldsymbol{w} \in \mathbb{R}^C$. In our problem, the rating vector is used as a weight vector to combine the shape responses of different models in the pool to estimate the final landmark positions. Similarly, the rating vector is also used to obtain the final visibility estimation. The training details of the ensemble of recommendation trees are provided in Sec. 3.2. Before we detail the training steps, we first review the conventional classification forest and discuss its limitations as a recommendation framework.

[3]Variations in landmark visibility are caused by both self-occlusion due to extreme view pose, and external-object occlusion.

3.1. Classification Forest for Model Recommendation

A classification forest can be adapted for model recommendation by training it to pick up the most suitable model given the model shape responses obtained from the model pool. The rating may be approximated by using the posterior classification probability. Alternatively, a hard model selection can be done by finding the maximum votes from all trees. In this section, we first give a brief review of classification forest and discuss its limitation as a model recommendation method.

Learning a conventional decision tree, *e.g.* the classification tree, typically starts from splitting the training set with optimal binary split function $h_\Theta(\boldsymbol{\phi}) \in \{0,1\}$, and repeats such process recursively until some criteria are reached. Typically the form of h is pre-defined, and for conventional classification tree, it could be represented by

$$h_\Theta(\boldsymbol{\phi}) = \begin{cases} 0 & \phi_k \le \tau \\ 1 & \phi_k > \tau \end{cases} . \tag{1}$$

A sample is channeled to the left child if $h_\Theta(\boldsymbol{\phi}) = 0$, and to the right if otherwise.

To choose the best split function parameters $\Theta_j = \{k_j, \tau_j\}$ for node j, a set of candidate features and thresholds are randomly selected and evaluated. The optimal one that maximises the information gain is chosen as the node parameters. Concretely, let denote the set of class labels as $\mathcal{L} = \{l^m\}_{m=1}^M$, where $l \in \{1, \ldots, C\}$ is the model label. Given the current split node j and its related sub training set $\mathcal{S}_j = \mathcal{F}_j \times \mathcal{L}_j \subset \mathcal{F} \times \mathcal{L}$, and let denote the two child nodes' sub training set by $\{\mathcal{S}_j^{(L)}, \mathcal{S}_j^{(R)}\}$, the information gain is defined as

$$I_j(\Theta_j) = H(\mathcal{S}_j) - \sum_{\theta\in\{L,R\}} \frac{|\mathcal{S}_j^{(\theta)}|}{|\mathcal{S}_j|} H(\mathcal{S}_j^{(\theta)}) \tag{2}$$

In a classification setting, $H(\mathcal{S}_j)$ denotes the cross entropy of sub training set $\mathcal{S}_j$, represented as

$$H(\mathcal{S}_j) = \sum_{c=1}^{C} -(p_j)_c \log(p_j)_c, \tag{3}$$

where p_j is the probability distribution calculated on $\mathcal{L}_j$.

Inference of classification forest can be achieved by feeding the feature $\boldsymbol{\phi}$ to all the trees and pass it downward according to the binary split function in each split node. Each tree would output a predicted label proposed by the resulted leaf node. We simply aggregate the decisions from all trees using a majority voting.

Despite that the classification forest is effective in handling high-dimensional features of shape responses, it is not

well-suited as a model recommendation framework for our problem, for the following reasons:

(1) The rating generated using the posterior class probability of a classification forest tends to simultaneously recommend non-related models far away from each other in the model space (*e.g.* left-profile and right-profile models), as depicted in Fig. 3(a). Such a poor recommendation will lead to erroneous landmark estimation. The poor recommendation is mainly caused by the splitting function of classification forest, in which the cross entropy formulation (Eq. (3)) does not naturally distinguish the different models. This drawback is resolved with our new splitting function. Fig. 3 shows an illustration that highlights the difference between the ratings generated by the classification tree and a recommendation tree. The rating vector generated by the recommendation tree tends to recommend models close in the shape space. We show in Sec. 4 that the proposed recommendation trees outperform the classification trees.

(2) The classification forest does not fully exploit the available supervision signal, namely landmark positions $\mathcal{Y}$ and visibility $\mathcal{V}$ in the training process.

3.2. Learning Ensemble of Model Recommendation Trees

We wish to address the shortcomings of classification forest through clustering similar faces naturally with a specially designed split mechanism. There are three main differences between the learning of a recommendation tree and a classification tree: (1) In a recommendation tree, the splitting of samples are directly related to their facial landmark locations instead of the expected model index as classification target. (2) To learn a recommendation tree, we need to provide the full supervision signal $\{\mathcal{F}, \mathcal{X}, \mathcal{Y}, \mathcal{V}\}$ rather than just $\mathcal{F} \times \mathcal{L}$, as in the classification forest. (3) A recommendation tree incorporates the rating vector $\boldsymbol{w}_j \in \mathbb{R}^C$ into the optimisation of each node j. All these differences are introduced aiming to cluster training samples with similar rating to the same tree node, and the similarity of rating is computed purely based on $\mathcal{X}$ and $\mathcal{Y}$.

More precisely, given current split node j and its related sub training set $\mathcal{S}_j = \{\mathcal{F}_j, \mathcal{X}_j, \mathcal{Y}_j, \mathcal{V}\}$, we define a new training objective as follows

$$I_j(\Theta_j, \boldsymbol{w}_j^{(\theta)}) = H(\mathcal{S}_j, \boldsymbol{w}_j) - \sum_{\theta \in \{L,R\}} \frac{|\mathcal{S}_j^{(\theta)}|}{|\mathcal{S}_j|} H(\mathcal{S}_j^{(\theta)}, \boldsymbol{w}_j^{(\theta)}), \tag{4}$$

where $H(\mathcal{S}_j, \boldsymbol{w}_j)$ is defined as

$$H(\mathcal{S}_j, \boldsymbol{w}_j) = \sum_{\substack{m,n: \\ \boldsymbol{x}^m \in \mathcal{S}_j, \\ \{m,n\} \in \mathcal{V}}} \|\boldsymbol{y}_n^m - \sum_{c=1}^{C} w_{jc} \boldsymbol{x}_{c,n}^m\|_2^2. \tag{5}$$

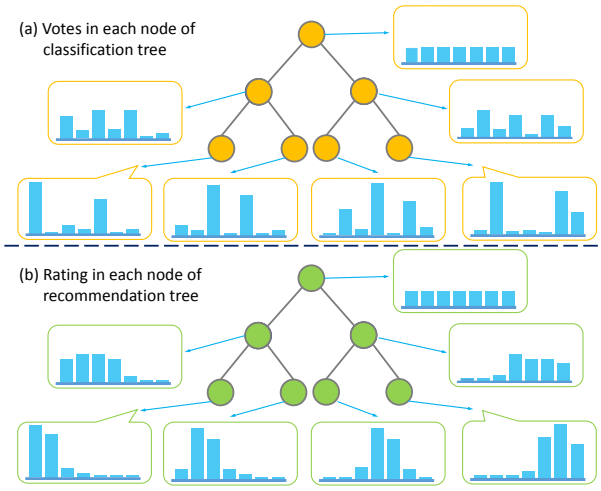

Figure 3. An illustration that highlights the differences between a conventional classification tree and a recommendation tree. To simplify the discussion, we assume the models in a pool are only yaw-specific without considering in-plane rotation or external-object occlusion. In each rating vector, the indices of rating are sorted based on the models' yaw angle. As can be observed, a classification tree neglects the potential relations between models, and may thus propose erroneous rating that advocates both left-profile-view and right-profile-view models. In contrast, the learning of a recommendation tree is directly supervised by facial landmark positions. Hence samples with similar facial shapes are naturally clustered, and hence similar models tend to share similar ratings. As a result, more robust and reasonable rating can be obtained.

Note that both $\mathcal{Y}$ and $\mathcal{V}$ are directly involved in the optimisation function, in contrast to the conventional classification forest. Importantly, the information gain is maximised not only based on Θ_j, but also the rating vectors in the parent and child nodes $\{\boldsymbol{w}_j, \boldsymbol{w}_j^{(L)}, \boldsymbol{w}_j^{(R)}\}$.

Optimisation with Eq. (4) is more complex than Eq. 2, which is used in the conventional classification forest. In the classification forest, optimisation is only performed over the split parameters Θ_j, where we just enumerate several candidate features and thresholds to find the optimal Θ_j. On the contrary, the optimisation in the recommendation forest needs to account for both the split parameters Θ_j and rating $\boldsymbol{w}_j^{(\theta)}(\Theta_j), \theta \in \{L, R\}$. In particular, for each candidate feature-threshold pair, we need to find the optimal $\boldsymbol{w}_j^{(\theta)}(\Theta_j), \theta \in \{L, R\}$. Assume that the Θ_j is given, the partition of $\mathcal{S}_j$ is determined, the optimal $\boldsymbol{w}_j^{(\theta)}(\Theta_j)$ can be obtained by

$$\begin{aligned} \boldsymbol{w}_j^{(\theta)}(\Theta_j) = \underset{\boldsymbol{w}}{\operatorname{argmin}} & \sum_{\substack{m,n: \\ \boldsymbol{x}^m \in \mathcal{S}_j^{(\theta)}, \\ \{m,n\} \in \mathcal{V}}} \|\boldsymbol{y}_n^m - \sum_{c=1}^{C} w_c \boldsymbol{x}_{c,n}^m\|_2^2 \\ s.t. \quad & \boldsymbol{w} \geq 0, \|\boldsymbol{w}\|_1 = 1, \end{aligned} \tag{6}$$

which can be efficiently solved by Quadratic Programming packages. Note that the regularity constraints over $\boldsymbol{w}$ in Eq. 6 is important to avoid unreasonable recommendation results and to enforce facial shape constraints.

3.3. Landmark Inference with EMRT

Despite the learning procedure of the proposed recommendation tree is more sophisticated than a traditional tree, the inference efficiency is identical to the conventional tree. Specifically, given an image in the testing stage, we first extract the shape responses $\tilde{\boldsymbol{x}}_{c,n}$ and subsequently compute the recommendation feature $\tilde{\boldsymbol{\phi}}$. The feature is fed to a recommendation tree and passed down from the root node to the leaf node according to the binary split function $h_{\Theta_j}(\tilde{\boldsymbol{\phi}})$ stored in each node. For each tree, rating $\boldsymbol{w}$ stored in the reached leaf node is used to estimate the predicted landmark locations $\tilde{\boldsymbol{y}}_n$ and visibility confidence $\tilde{v}_n$,

$$\tilde{\boldsymbol{y}}_n = \sum_{c=1}^{C} w_c \tilde{\boldsymbol{x}}_{c,n}, \qquad \tilde{v}_n = \sum_{c:(\boldsymbol{b}_c)_n=1} w_c \tilde{\phi}_{c,n}, \tag{7}$$

where $\tilde{\phi}_{c,n}$ is retrieved from the F-dimensional $\tilde{\boldsymbol{\phi}}$.

To obtain the final estimation from different trees, we average all the predictions (landmark locations and visibility confidence) weighted by the number of training samples in the reached leaf node of each tree. Visibility of a landmark is obtained via comparing the estimated confidence with a threshold γ.

4. Experiments

Datasets. Following the objective of this study of aligning arbitrary-view face up to profile view, we choose the following datasets with extremely large pose variations.

AFLW dataset [14]: This dataset was collected from Flickr hence exhibits in-the-wild setting. The original dataset contains 25,993 images, of which 27.48% have yaw angle larger than 45° and 3.83% larger than 90°. The large angle range and in-the-wild setting is challenging for common methods without specifically designed mechanism. Each face is annotated with at most 21 landmarks. Similar to [26, 1], we discard two ear landmarks to enforce more stable face shape. Non-visible landmark annotation, including both self-occlusion and external-object-occlusion, is not provided. To our knowledge, our study is the first attempt to train and evaluate on this challenging dataset without removing non-frontal faces.

MultiPIE dataset [20]: This dataset was originally collected in a constrained lab environment with a setting of multiple poses, illuminance and facial expressions. A subset of 6152 faces are defined [20] and labelled with either 39 (profile-view) or 68 (frontal-view) facial landmarks. Zhu *et al*. [32] further annotate another 400 profile-view faces with 39 landmarks. In our experiments, these two subsets of data are merged and cleaned to form a 6547-image dataset.

AFW dataset [32]: This original dataset contains 468 faces with in-the-wild setting. Each face is annotated with at most 6 landmarks. Non-visible landmark annotation is not provided similar to AFLW. We note that a subset of AFW, totally 337 faces, were selected and re-annotated with 68 landmarks by [18]. These 337 faces are frontally biased, hence we still employ the original 468 faces and the associated annotations.

Evaluation metric. We apply percentage of mean squared error normalised by the upright distance between the highest eyebrow landmark and the chin landmark as the evaluation metric. Errors are calculated on all annotated landmarks, excluding non-annotated or invisible landmarks, since no ground truths are available. We do not normalise error by inter-ocular distance because one of the two eyes is invisible for many profile view cases.

4.1. Comparison with Existing Methods

Comparison with [32]. To our knowledge, [32] is one of the most recent works that is capable of handling profile-view and frontal-view faces simultaneously. Model of [32] is trained on only 900 images of MultiPIE dataset, which is a rather small fraction of the original dataset. Since the training of a discriminative regression model requires more training samples, we employ the full MultiPIE dataset and perform evaluations by averaging performance on 5-fold cross validation. For a fair comparison, we re-train the model of [32] in the same way using their released codes. Note that each person might contribute multiple facial images. We ensure that the images from the same person only appear within the same fold.

Similar to [32], we further evaluate the models trained on the full MultiPIE set on the **full** AFW dataset. As in [32], we use a trained linear mapping to obtain the 6-point estimate by each method. Since [32] predefines the mixture of trees, it is infeasible to train the model of [32] with AFLW landmarks protocol. Hence we do not compare with [32] on the AFLW dataset.

Approach	MultiPIE	AFW
[32]	2.74	4.68
Ours	**2.12**	**3.55**

Table 2. Mean error comparison with [32] on two datasets (also applied in [32]).

Results of the comparison are reported in Table 2. We also compare the relative cumulative error distribution (CED) curve for the two methods on MultiPIE and AFW datasets in Fig. 4. Based on the results, it is observed that our framework outperforms [32] on both the datasets applied in their work. More significant improvement is ob-

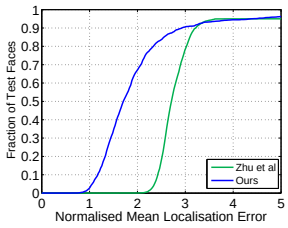

(a) Results from MultiPIE.

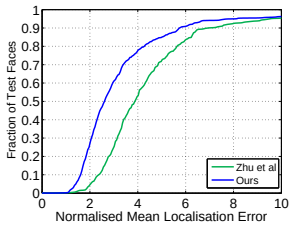

(b) Results from AFW.

Figure 4. CED curve compared with [32] on MultiPIE and AFW.

served on AFW dataset with in-the-wild setting. The results suggest the effectiveness of using the discriminative method along with the proposed recommendation framework for handling unconstrained settings. We provide qualitative examples in Fig .5 comparing against [32].

Comparison with [24, 17, 13]. The methods proposed in these three studies are representative discriminative regression approaches. They have achieved state-of-the-art results on commonly used frontal-view-biased datasets. In this part we compare our proposed recommendation framework with these regression methods to show the benefits of adopting a model recommendation framework. Since the training codes of [24, 17] are not publicly available, we evaluate all these methods based on our re-implementation. Our re-implementation achieves comparable performance on common benchmarks (*e.g.* LFPW [1], Helen [15]) with theirs reported in the literature.

We evaluate our approach and the three baseline methods on the **full** AFLW dataset since it is more challenging than the MultiPIE dataset. Since AFLW also does not provide a train-test set partition, we perform a 5-fold cross-validation and report the averaged performance across all the folds.

Approach	Mean error over AFLW dataset
SDM [24]	6.98
LBF [17]	7.47
ERT [13]	7.70
Ours	**4.01**

Table 3. Mean error comparison with [24, 17, 13] on the full AFLW dataset.

Note that we only train the discriminative regression methods with the frontal face subset. Training the baseline methods using the full training set is infeasible since AFLW contains varying point sets and these methods cannot handle this issue. Simply filling the missing points with transformed mean shape is possible, but training the baseline methods with this augmented point sets lead to worse performance than using only the frontal face subset.

The comparative results are summarised in Table 3. The results suggest the importance of deploying multiple models together with a model recommendation framework to handle arbitrary views. Qualitative results by our method are shown in Fig. 6.

4.2. Further Analysis

To validate the effectiveness of the proposed model recommendation trees, we further compare our approach with some variants derived from our framework and alternative strategies:

(a) No recommendation at all, *i.e.* always choosing the frontal model as recommended model. This variant is equivalent to SDM [24];

(b) Prior head pose estimation before face alignment. This method is conceptually similar to the conditional regression forest [8];

(c) Optimising landmarks locations based on multi-response map. This method is conceptually similar to the study on consensus by regressors [27];

(d) Classification forest choosing the model with top voting;

(e) Classification forest using posterior class probability as rating.

Note that all the baselines exploits the same SDM models as those used in our approach to ensure a fair comparison in the model recommendation process. We compare all the methods on the challenging full AFLW dataset and the training setting is similar to that described in Sec. 4.1. We adopt three evaluation metrics, namely mean normalised error, visibility detection accuracy, and the corresponding average precision. Visibility detection accuracy refers to the percentage of landmarks predicted with correct visibility.

Results are provided in Table 4. Precision-Recall curve of visibility evaluation is depicted in Fig. 7. Following [3], we treat invisible landmarks as 'positive'. The results suggest the effectiveness of the the proposed model recommendation trees. It can be observed that an application of approaches with prior head pose estimation [8] or occlusion handling [27] only gains marginal improvement compared with the frontal-face SDM on the challenging arbitrary-view face alignment task. Moreover, a naive application of classification forest (both choosing top voted model or using posterior class probability as rating) cannot achieve the desired results.

5. Conclusion

We have presented a novel model recommendation framework to simultaneously handle face alignment of arbitrary-view faces up to profile view. The framework can

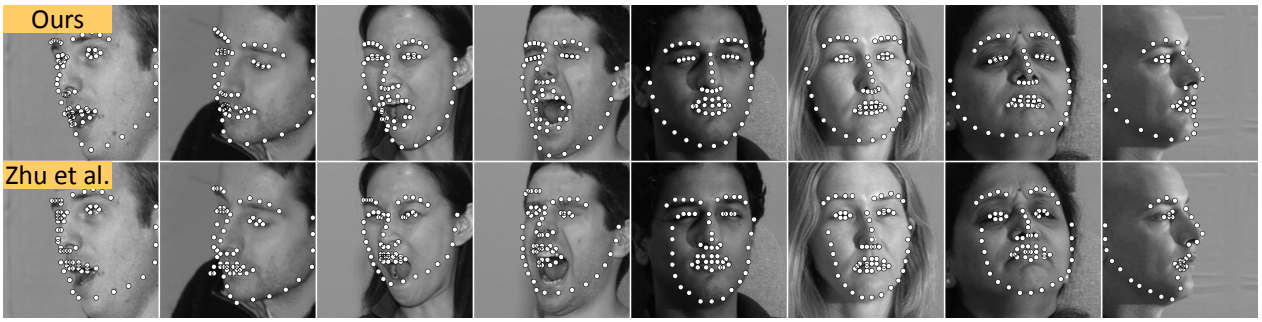

Figure 5. Comparison of representative results between our method (first row) and [32] (second row).

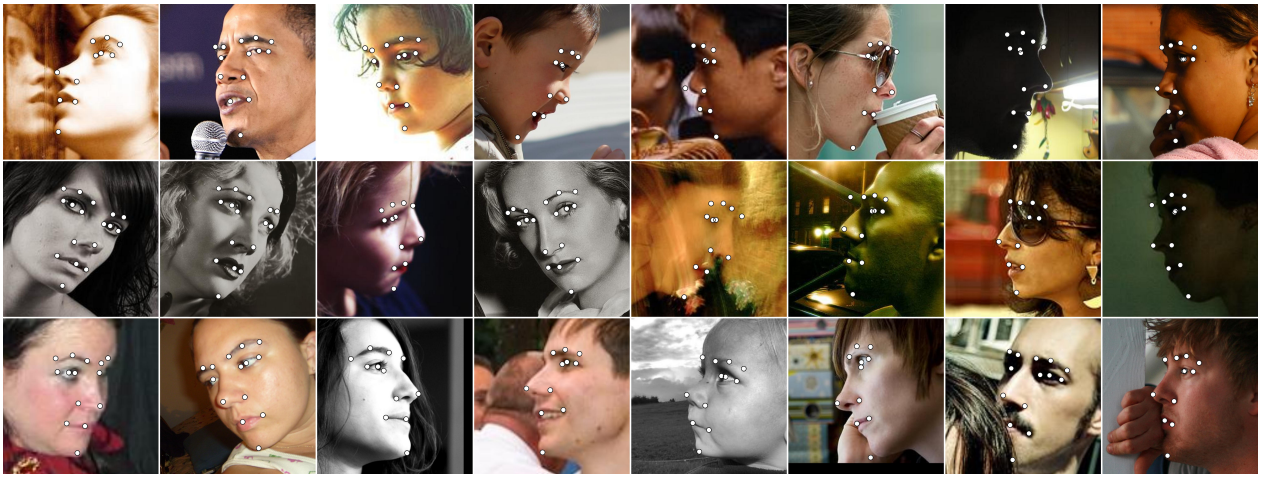

Figure 6. Illustrative results by our method on AFLW dataset. It is worth pointing that all faces have extreme view pose (up to profile view or with in-plane rotation).

	Mean error on AFLW	Visibility detection accuracy (%)	Visibility AP (%)
(a)	6.98	83.20	42.35
(b)	6.61	83.33	41.78
(c)	6.65	87.89	47.08
(d)	9.24	65.26	25.75
(e)	7.62	87.65	45.79
Ours	**4.01**	**91.18**	**49.35**

Table 4. Comparison of model recommendation between the proposed recommendation trees and various baseline methods.

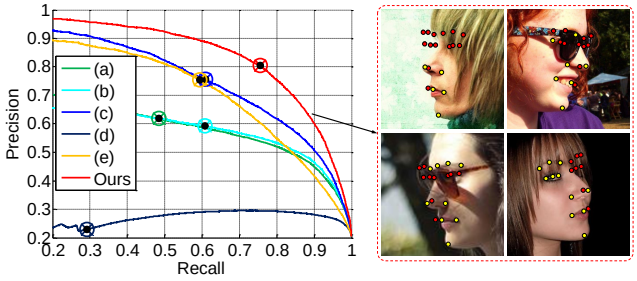

Figure 7. Visibility detection Precision-Recall curve compared to various baseline methods, with representative samples. The black dot on the curve denotes the visibility threshold. Threshold for each method is chosen to maximise visibility detection accuracy on the validation set. Colors of the annotated landmarks indicate our estimate of visibility (yellow for visible and red for invisible). Best viewed in color.

flexibly work with different existing discriminative regression based methods. It is capable of accommodating different landmark protocols with varying point sets and inferring visibility of landmarks. Extensive experiments on the full set of AFLW, MultiPIE, and AFW demonstrate the effectiveness of our approach over existing multimodal face alignment methods.

References

[1] P. N. Belhumeur, D. W. Jacobs, D. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. In *CVPR*, pages 545–552, 2011. 1, 2, 3, 6, 7

[2] L. Breiman. Random forests. *Machine learning*, 45(1):5–32,

2001. 2

[3] X. P. Burgos-Artizzu, P. Perona, and P. Dollár. Robust face landmark estimation under occlusion. In *ICCV*, pages 1513–1520, 2013. 2, 3, 7

[4] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. *IJCV*, 107(2):177–190, 2014. 2

[5] C. Chen, A. Dantcheva, and A. Ross. Automatic facial makeup detection with application in face recognition. In *ICB*, pages 1–8, 2013. 1

[6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *TPAMI*, 23(6):681–685, 2001. 2

[7] T. F. Cootes, M. C. Ionita, C. Lindner, and P. Sauer. Robust and accurate shape model fitting using random forest regression voting. In *Computer Vision–ECCV 2012*, pages 278–291. Springer, 2012. 2

[8] M. Dantone, J. Gall, G. Fanelli, and L. Van Gool. Real-time facial feature detection using conditional regression forests. In *CVPR*, pages 2578–2585, 2012. 2, 3, 7

[9] A. Datta, R. Feris, and D. Vaquero. Hierarchical ranking of facial attributes. In *AFGR*, pages 36–42. IEEE, 2011. 1

[10] P. Dou, Y. Wu, S. K. Shah, and I. A. Kakadiaris. Benchmarking 3d pose estimation for face recognition. In *Pattern Recognition (ICPR), 2014 22nd International Conference on*, pages 190–195. IEEE, 2014. 1

[11] S. Du, Y. Tao, and A. M. Martinez. Compound facial expressions of emotion. *Proceedings of the National Academy of Sciences*, 111(15):E1454–E1462, 2014. 1

[12] G. Hermosilla, P. Loncomilla, and J. Ruiz-del Solar. Thermal face recognition using local interest points and descriptors for hri applications. In *RoboCup 2010: Robot Soccer World Cup XIV*, pages 25–35. Springer, 2011. 1

[13] V. Kazemi and S. Josephine. One millisecond face alignment with an ensemble of regression trees. In *CVPR*, 2014. 1, 2, 7

[14] M. Kostinger, P. Wohlhart, P. M. Roth, and H. Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *ICCVW*, pages 2144–2151, 2011. 2, 6

[15] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang. Interactive facial feature localization. In *ECCV*, pages 679–692, 2012. 1, 7

[16] P. Matikainen, R. Sukthankar, and M. Hebert. Model recommendation for action recognition. In *CVPR*, pages 2256–2263, 2012. 3

[17] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment at 3000 fps via regressing local binary features. In *CVPR*, 2014. 1, 2, 7

[18] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCVW*, pages 397–403, 2013. 3, 6

[19] X. Shen, Z. Lin, J. Brandt, and Y. Wu. Detecting and aligning faces by image retrieval. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3460–3467. IEEE, 2013. 2

[20] T. Sim, S. Baker, and M. Bsat. The cmu pose, illumination, and expression database. *TPAMI*, 25(12):1615–1618, 2003. 2, 6

[21] Y. Sun, Y. Chen, X. Wang, and X. Tang. Deep learning face representation by joint identification-verification. In *Advances in Neural Information Processing Systems*, pages 1988–1996, 2014. 1

[22] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. DeepFace: Closing the gap to human-level performance in face verification. In *CVPR*, 2014. 1

[23] J. Xing, Z. Niu, J. Huang, W. Hu, and S. Yan. Towards multi-view and partially-occluded face alignment. In *CVPR*, pages 1829–1836, 2014. 1, 3

[24] X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *CVPR*, pages 532–539, 2013. 1, 2, 7

[25] C.-Y. Yang, S. Liu, and M.-H. Yang. Structured face hallucination. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1099–1106. IEEE, 2013. 1

[26] H. Yang and I. Patras. Sieving regression forest votes for facial feature detection in the wild. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1936–1943. IEEE, 2013. 2, 6

[27] X. Yu, Z. Lin, J. Brandt, and D. N. Metaxas. Consensus of regression for occlusion-robust facial feature localization. In *ECCV*, pages 105–118, 2014. 1, 2, 3, 7

[28] G. Zen, E. Sangineto, E. Ricci, and N. Sebe. Unsupervised domain adaptation for personalized facial emotion recognition. In *Proceedings of the 16th International Conference on Multimodal Interaction*, pages 128–135. ACM, 2014. 1

[29] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, pages 94–108, 2014. 2

[30] X. Zhao, T.-K. Kim, and W. Luo. Unified face analysis by iterative multi-output random forests. In *CVPR*, pages 1765–1772, 2014. 2, 3

[31] Y. Zhou, W. Zhang, X. Tang, and H. Shum. A Bayesian mixture model for multi-view face alignment. In *CVPR*, volume 2, pages 741–746. IEEE, 2005. 1, 3

[32] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *CVPR*, pages 2879–2886, 2012. 1, 2, 6, 7, 8